

**IN THE UNITED STATES
PATENT AND TRADEMARK OFFICE**

Patent Application

Inventors: Sachin Garg et al.

Serial No.: 10/662728

Conf. No.: 1515

Filing Date: 9/15/2003

Art Unit: 2143

Examiner: Mark Fearer

Docket No.: 630-045US

Title: Congestion Management in Telecommunications Networks

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

APPEAL BRIEF UNDER 37 CFR 41.67

Pursuant to 37 CFR 41.67, this brief is filed in support of the appeal in this application.

TABLE OF CONTENTS

| | |
|--|----|
| REAL PARTY IN INTEREST | 4 |
| RELATED APPEALS AND INTERFERENCES | 5 |
| STATUS OF CLAIMS..... | 6 |
| STATUS OF AMENDMENTS | 7 |
| SUMMARY OF THE CLAIMED SUBJECT MATTER..... | 8 |
| GROUNDS OF OBJECTION AND REJECTION TO BE REVIEWED ON APPEAL | |
| Ground 1: 35 U.S.C. 103 Rejection of Claims 1-2, 4-6, and 8-10 | 14 |
| Ground 2: 35 U.S.C. 103 Rejection of Claims 3 and 7..... | 14 |
| ARGUMENTS | |
| Ground 1: 35 U.S.C. 103 Rejection of Claims 1-2, 4-6, and 8-10 | 15 |
| Ground 2: 35 U.S.C. 103 Rejection of Claims 3 and 7..... | 16 |
| CONCLUSION | 17 |
| CLAIMS APPENDIX | 18 |
| EVIDENCE APPENDIX | 20 |
| RELATED PROCEEDINGS APPENDIX..... | 21 |

REAL PARTY IN INTEREST

The real party of interest in this application is the assignee of this application: Avaya Technology Corp., of Basking Ridge, NJ.

RELATED APPEALS AND INTERFERENCES

U.S. patent application Serial No. 10/662,724, filed 09/15/2003 (Attorney Docket: 630-044us) is related to this application. An appeal in that case is currently pending and awaiting review.

STATUS OF CLAIMS

Claims 1-10 stand rejected and are being appealed.

STATUS OF AMENDMENTS

All amendments have been entered.

SUMMARY OF THE CLAIMED SUBJECT MATTER

A typical computer network comprises a plurality of nodes and links and is depicted in Figure 1.




Figure 1 – A Typical Computer Network

Each link carries information from one node to another and can be, for example, an electrical cable, an optical cable, or a wireless relay. Each node switches information from one link to another and can be, for example, a switch, a router, or an access point. Applicant's Specification [0008]

Most computer networks transmit information in discrete chunks called "protocol data units." Frames, packets, and datagrams are typical protocol data units. Applicant's Specification [0003] Protocol data units are injected into a network at one node and are passed from node to node, in bucket-brigade fashion, until they arrive at their destination. This is illustrated in Figure 2. Applicant's Specification [0002]




Figure 2 – Protocol Data Units Being Passed Bucket-Brigade Fashion

In some cases, a protocol data unit might spend a relatively short amount of time in a node before it is processed and transmitted on an output link. In other cases, a protocol data unit might spend a relatively long time. Applicant's Specification [0004]

One reason why a protocol data unit might spend a long time in a network node is because the output link on which the protocol data unit is to be transmitted is temporarily unavailable. Another reason why a protocol data unit might spend a long time in a network node is because a large number of protocol data units arrive at the node faster than the node can process and output them. Applicant's Specification [0005]

To accommodate this, a network node typically stores or "queues" a protocol data unit until it is transmitted. Sometimes, the protocol data units are stored in an "input queue" and sometimes the protocol data units are stored in an "output queue." An input queue might be employed when protocol data units arrive at the network node (in the short run) more quickly than they can be processed. An output queue might be employed when protocol data units arrive and are processed (in the short run) more quickly than they can be transmitted on the output link. Applicant's Specification [0006]

A queue has a finite capacity, and, therefore, it can fill up with protocol data units. When a queue is filled, the attempted addition of protocol data units to the queue causes

the queue to “overflow” with the result that the newly arrived protocol data units are discarded or “dropped.” Dropped protocol units are lost forever and do not leave the network node. Applicant’s Specification [0007]

A network node that comprises a queue that is dropping protocol data units is called “congested.” For the purposes of this specification, a “**congestible node**” is defined as a network node (*e.g.* a switch, router, access point, *etc.*) that is susceptible to dropping protocol data units. Applicant’s Specification [0008]

The loss of a protocol data unit is detrimental to the user of the protocol data unit, but the loss of any one protocol data unit does not have the same degree of impact as every other protocol data unit. In other words, the loss of some protocol data units is more injurious than the loss of some others. Applicant’s Specification [0009]

When a congestible node is congested, or close to becoming congested, it can be prudent for the node to intentionally and proactively drop one or more protocol data units whose loss will be less consequential than to allow arriving protocol data units to overflow and be dropped and whose loss might be more consequential. To accomplish this, the node can employ an intelligent congestion management algorithm to decide:

- which protocol data units to drop,
- how many protocol data units to drop, and
- when to drop those protocol data units,

in order to:

- reduce injury to the affected communications, and
- lessen the likelihood of congestion in the congestible node.

Applicant’s Specification [0010] How intelligent congestion management algorithms are designed is well known in the prior art and is not germane to the present invention.

Before intelligent congestion management algorithms were invented, there were thousands of congestible nodes in use. Today — after the development of intelligent congestion management algorithms — there are still thousands of congestible nodes in use. Why? Because intelligent congestion management algorithms cannot be retrofitted into most congestible nodes.

The solution to the problem of the existence of thousands of old congestible nodes without intelligent congestion management is — at least according to the prior art — to replace the old nodes with new nodes that do incorporate intelligent congestion management. The inventors of the present invention recognized that this answer is prohibitively expensive and in many cases suggests the replacement of nodes that are otherwise perfectly-good. Clearly, a better solution was needed than to swap out thousands of legacy nodes. Furthermore, the proliferation of WiFi, Bluetooth, and Zigbee networks is fueling the need for inexpensive “lightweight” nodes that do not have the computing power to implement intelligent congestion management.

In response, the applicants invented a device that:

1. can perform intelligent congestion management for a congestible node that cannot perform it for itself, and
2. has a queue which effectively supplements the storage capacity of the queue in the congestible node, and
3. can be added to a computer network without any reconfiguration of the network, and
4. can be added to the computer network without any reconfiguration of the congestible node.

The device — called a “protocol-data-unit excisor” — is inserted into the link carrying protocol data units to a congestible node. This is illustrated in Figure 3.

The protocol-data-unit excisor comprises a queue that can store protocol data units when the congestible node is not yet ready to receive them, and the protocol-data-unit excisor transmits protocol data units to the congestible node from its queue when the congestible node indicates that it is ready to receive them. The congestible node signals the protocol-data-unit excisor that it is or is not ready to receive protocol data units through traditional flow control signals.

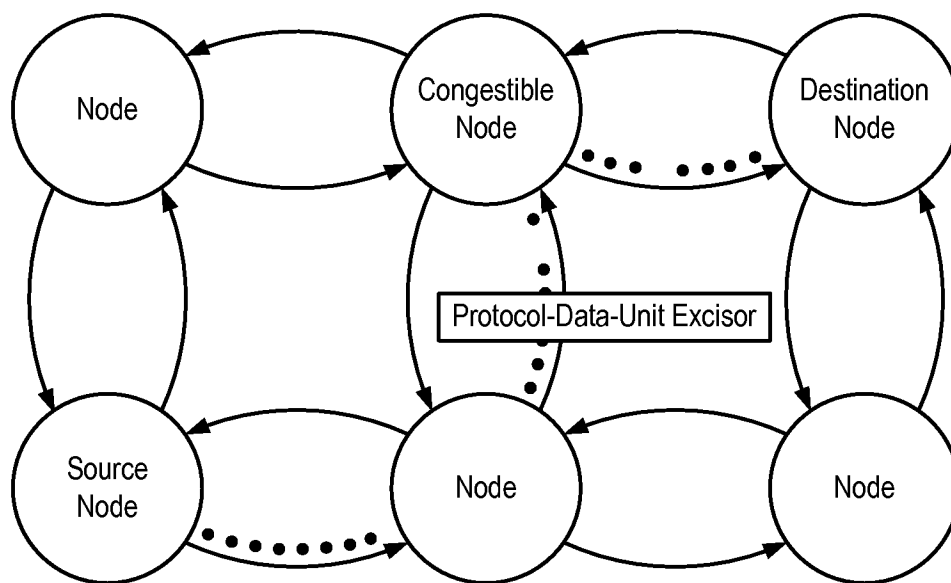


Figure 3 – Protocol-Data-Unit Excisor Inserted Into Link To Congestible Node

The protocol-data-unit excisor also performs intelligent congestion management on its own queue, which effectively eliminates the need for the congestible node to do it for itself.

There are two pending independent claims at issue in this appeal.

Claim 1 recites:

1. A method comprising:

receiving a first plurality of protocol data units at a first input, wherein all of said first plurality of protocol data units are en route to a first congestible node;

maintaining at a protocol-data-unit excisor a first queue for said first plurality of protocol data units;

receiving at said protocol-data-unit excisor a flow control signal that indicates whether said first congestible node is ready to receive one or more of said protocol data units from said first queue; and

selectively dropping, at said protocol-data-unit excisor, one or more of said protocol data units based on a first metric of said first queue.

Claim 1 is supported in the specification at **Paragraphs [0047]-[0064] and Figures 6 and 7.**

Independent claim 6 is the apparatus equivalent of claim 1, which recites:

6. A protocol-data-unit excisor comprising:

- a first input for receiving a first plurality of protocol data units, wherein all of said first plurality of protocol data units are *en route* to a first congestible node;
- a first queue for storing said first plurality of protocol data units;
- a first receiver for receiving a flow control signal that indicates whether said first congestible node is ready to receive one or more of said protocol data units from said first queue; and
- a processor for selectively dropping one or more of said protocol data units based on a metric of said first queue.

Claim 6 is supported in the specification at **Paragraphs [0024]-[0046] and Box 302 in Figure 3 and all of Figure 4.**

GROUND OF OBJECTION AND REJECTION TO BE REVIEWED ON APPEAL

Ground 1: 35 U.S.C. 103 Rejection of Claims 1-2, 4-6, and 8-10

Claims 1-2, 4-6, and 8-10 have been rejected under 35 U.S.C. 103(a) as being unpatentable over S. Miller et al., U.S. Patent 6,650,640 B1 (hereinafter "Miller") in view of B. Erimli et al., U.S. Patent 6,405,258 B1 (hereinafter "Erimli").

Ground 2: 35 U.S.C. 103 Rejection of Claims 3 and 7

Claims 3 and 7 have been rejected under 35 U.S.C. 103(a) as being unpatentable over S. Miller et al., U.S. Patent 6,650,640 B1 (hereinafter "Miller") in view of B. Erimli et al., U.S. Patent 6,405,258 B1 (hereinafter "Erimli") and further in view of S. Yu, U.S. Patent 7,031,341.

ARGUMENTS

Ground 1: 35 U.S.C. 103 Rejection of Claims 1-2, 4-6, and 8-10

Claims 1-2, 4-6, and 8-10 have been rejected under 35 U.S.C. 103(a) as being unpatentable over S. Miller et al., U.S. Patent 6,650,640 B1 (hereinafter “Miller”) in view of B. Erimli et al., U.S. Patent 6,405,258 B1 (hereinafter “Erimli”). The applicants respectfully traverse.

Claim 1 recites:

1. A method comprising:
receiving a first plurality of protocol data units ***at a first input, wherein all of said first plurality of protocol data units are en route to a first congestible node;***
maintaining at a protocol-data-unit excisor a first queue for said first plurality of protocol data units;
receiving at said protocol-data-unit excisor a flow control signal that indicates whether said first congestible node is ready to receive one or more of said protocol data units from said first queue; and
selectively dropping, at said protocol-data-unit excisor, one or more of said protocol data units based on a first metric of said first queue.
(emphasis supplied)

Nowhere Miller nor Erimli, alone or in combination, teach or suggest, what claim 1 recites — namely, receiving a first plurality of protocol data units at a first input, wherein all of the protocol data units are en route to a first congestible node. In other words, all of the protocol data units that arrive at one input are destined for one congestible node — not one of two or three nodes — but exactly one node. The purpose of this limitation is to exclude protocol-data-unit excisors — that perform switching (such as those in Miller) from the scope of the claim.

The Office action does not acknowledge this limitation nor does it provide a reference into Miller or Erimli to substantiate the rejection. And a careful reading of the references shows that they do not anticipate this limitation.

For this reason, the applicants respectfully submit that the rejection of claim 1 is traversed.

Because claims 2 and 4-5 depend on claim 1, the applicants respectfully submit that the rejection of them is also traversed.

Claim 6 recites:

6. (previously presented) A protocol-data-unit excisor comprising:
a first input for receiving a first plurality of protocol data units, wherein all of said first plurality of protocol data units are en route to a first congestible node;
a first queue for storing said first plurality of protocol data units;
a first receiver for receiving a flow control signal that indicates whether said first congestible node is ready to receive one or more of said protocol data units from said first queue; and
a processor for selectively dropping one or more of said protocol data units based on a metric of said first queue.
(emphasis supplied)

For essentially the same reasons as those given with respect to claim 1, the applicants respectfully submit that the rejection of it is traversed.

Because claims 8-10 depend on claim 6, the applicants respectfully submit that the rejection of them is also traversed.

Ground 2: 35 U.S.C. 103 Rejection of Claims 3 and 7

Claims 3 and 7 have been rejected under 35 U.S.C. 103(a) as being unpatentable over S. Miller et al., U.S. Patent 6,650,640 B1 (hereinafter "Miller") in view of B. Erimli et al., U.S. Patent 6,405,258 B1 (hereinafter "Erimli") and further in view of S. Yu, U.S. Patent 7,031,341 (hereinafter "Yu").

Because claim 3 depends on Yu fails to cure the deficiency of Miller and Erimli with respect to claim 1, the applicants respectfully submit that the rejection of claim 3 is traversed.

Because claim 7 depends on Yu fails to cure the deficiency of Miller and Erimli with respect to claim 6, the applicants respectfully submit that the rejection of claim 7 is traversed.

CONCLUSION

The applicants have demonstrated that the logic underlying the Office's rejection is untenable, and, therefore, that the rejection is not sustainable. For this reason, the applicants respectfully request the Board of Appeals to reverse the decision of the Examiner as provided for in 37 C.F.R. 41.50(a).

Respectfully,
Sachin Garg et al.

By **/Jason Paul DeMont/**
Jason Paul DeMont
Reg. No. 35,793
Attorney for Applicants
732-687-7990

DeMont & Breyer, L.L.C.
Suite 250
100 Commons Way
Holmdel, NJ 07733
United States of America

Claims Appendix

1. (previously presented) A method comprising:

receiving a first plurality of protocol data units at a first input, wherein all of said first plurality of protocol data units are en route to a first congestible node;

maintaining at a protocol-data-unit excisor a first queue for said first plurality of protocol data units;

receiving at said protocol-data-unit excisor a flow control signal that indicates whether said first congestible node is ready to receive one or more of said protocol data units from said first queue; and

selectively dropping, at said protocol-data-unit excisor, one or more of said protocol data units based on a first metric of said first queue.

2. (previously presented) The method of claim 1 wherein said protocol-data-unit excisor decides whether to drop a protocol data unit based on Random Early Detection.

3. (previously presented) The method of claim 1 wherein said indication is conveyed using back-pressure flow control.

4. (previously presented) The method of claim 1 wherein said indication is conveyed using the Pause frame procedure of IEEE 802.3.

5. (previously presented) The method of claim 1 further comprising:

receiving a second plurality of protocol data units at a second input, wherein all of said second plurality of protocol data units are en route to a second congestible node;

maintaining at said protocol-data-unit excisor a second queue for said for said second plurality of protocol data units;

receiving at said protocol-data-unit excisor a flow control signal that indicates whether said second congestible node is ready to receive one or more of said protocol data units from said second queue; and

selectively dropping, at said protocol-data-unit excisor, one or more of said protocol data units based on a second metric of said second queue.

6. (previously presented) A protocol-data-unit excisor comprising:
a first input for receiving a first plurality of protocol data units, wherein all of said first plurality of protocol data units are *en route* to a first congestible node;
a first queue for storing said first plurality of protocol data units;
a first receiver for receiving a flow control signal that indicates whether said first congestible node is ready to receive one or more of said protocol data units from said first queue; and
a processor for selectively dropping one or more of said protocol data units based on a metric of said first queue.

7. (previously presented) The protocol-data-unit excisor of claim 6 wherein said indication is conveyed using back-pressure flow control.

8. (previously presented) The protocol-data-unit excisor of claim 6 wherein said indication is conveyed using the Pause frame procedure of IEEE 802.3.

9. (previously presented) The protocol-data-unit excisor of claim 6 wherein said protocol-data-unit excisor decides whether to drop a protocol data unit based on Random Early Detection.

10. (previously presented) The protocol-data-unit excisor of claim 6 further comprising:

a second input for receiving a second plurality of protocol data units, wherein all of said second plurality of protocol data units are en route to a second congestible node;
a second queue for storing said second plurality of protocol data units; and
a second receiver for receiving a flow control signal that indicates whether said second congestible node is ready to receive one or more of said protocol data units from said second queue;
wherein said processor is also for selectively dropping one or more of said protocol data units based on a metric of said second queue.

Evidence Appendix

There is no evidence submitted pursuant to 37 CFR §§ 1.130, 1.131, or 1.132.

Related Proceedings Appendix

U.S. patent application Serial No. 10/662,724, filed 09/15/2003 (Attorney Docket: 630-044us) is related to this application. An appeal in that case is currently pending and awaiting review.